




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 6 : H04M 3/50, H04L 12/58, H04M 3/42		A1	(11) International Publication Number: WO 98/13993
			(43) International Publication Date: 2 April 1998 (02.04.98)

(21) International Application Number: PCT/GB97/02608	(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, US, UZ, VN, YU, ZW, ARIPO patent (GH, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).
(22) International Filing Date: 25 September 1997 (25.09.97)	
(30) Priority Data: 9619958.3 25 September 1996 (25.09.96) GB 9707712.7 16 April 1997 (16.04.97) GB	
(71) Applicant (for all designated States except US): BRITISH TELECOMMUNICATIONS PUBLIC LIMITED COMPANY [GB/GB]; 81 Newgate Street, London EC1A 7AJ (GB).	
(72) Inventor; and	Published
(75) Inventor/Applicant (for US only): GIBSON, David, Lynton [GB/GB]; "Firleigh", Struggler's Lane, Witnesham, Ipswich, Suffolk IP6 9HS (GB).	With international search report.
(74) Agent: DUTTON, Erica, Lindley, Graham; BT Group Legal Services, Intellectual Property Dept., Holborn Centre, 8th floor, 120 Holborn, London EC1N 2TE (GB).	

(54) Title: APPARATUS FOR COMMUNICATIONS SERVICE PROVISION



(57) Abstract

A computer applications platform (150) provides an interface between telecommunications networks, such as the Public Switched Telecommunications Network (160) in the UK, and data networks such as the Internet (155). Applications which run on the platform (150) can be used to call up resources, for instance for the delivery of recorded messages. Because the applications platform (150) bridges data and telecommunications networks, the services provided can be via either type of network. Therefore a user of the telecommunications network can trigger a fax or electronic mail message to be sent to a user on a data network. The service being requested by a user can be identified by the applications platform from a data field in an incoming data message, or by the direct dialled digits associated with an incoming telephony line.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Switzerland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		

APPARATUS FOR COMMUNICATIONS SERVICE PROVISION

The present invention relates to the provision of communications services and finds 5 particular application in the integration of different service types, such as voice telephony and messaging.

In conducting their everyday work, people need to communicate with each other. The nature of the communication will often reflect the nature of what it is wished 10 to impart and the facilities available to both the sender and the recipient. Thus, at present, a typical office will contain a facsimile (FAX) machine, with telephones and personal computers (PCs) at each desk so that each employee has these facilities at their disposal on a personal or shared basis.

15 These facilities are generally not well integrated - for example, if a user keeps a directory of regularly used telephone or FAX numbers on his PC he will typically have to re-enter these numbers on the telephone keypad when he wants to contact people by telephone or FAX. Indeed, having identified the person he needs to contact, the user is often uninterested in knowing the number to dial, he merely 20 wishes to be put in touch. This lack of integration is more apparent should he be away from his office and wish to access any received messages. In this circumstance he will generally have to ring different numbers for each messaging service, using different terminal equipment for each service which, on some occasions, may not be readily available.

25

There is a requirement, therefore, for a system which can provide improved access to both telephony and other services, such as messaging.

It is known to provide an interface between a computer network and a Public 30 Switched Telephone Network (PSTN) for this purpose. Apparatus which provides such an interface is a speech platform, such as that used by the present applicant to provide the "Callminder" Service in the United Kingdom.

The "Callminder" service is a network-based answering machine service. It records voice messages for a user who is not there to receive an incoming call. The speech platform which supports it clearly has to have substantial capacity to deal with the volume of traffic generated by the PSTN. However, there are other 5 services which can be supported from a speech platform.

According to the present invention, there is provided an applications platform for providing communications services, the platform comprising:

- 10 i) a connection channel interface, for providing access to two or more connection channels to a network for carrying voice signals;
- ii) a data network interface, for providing access to a network for carrying data;
- 15 iii) access to resources for use in providing communications services, including at least one speech-related resource from the group comprising voice recognition, recordal of incoming sound signals and transmission of outgoing sound signals;
- iv) means for receiving service requests at the connection channel interface;
- 20 v) interpretation means for use in relating a service request received at the connection channel interface to a computer-based application for use in provision of that service; and
- 25 vi) initiating means for identifying and initiating the running of one or more computer-based applications in response to service requests received at the connection channel interface, said application(s) being identified by reference to the interpretation means, which one or more computer-based application(s) is adapted to call on at least one of said resources to provide a communications
- 30 service in response to a received service request,
- such that the applications platform is able to respond to a received service request by identifying one or more computer-based applications for use in provision of that service and initiating the running of the identified application(s).

The word "application" is used in this context in its computer-related meaning and indicates computer software for providing and controlling processes such as communications services.

5

The network for carrying voice signals and the network for carrying data may be provided by the same network infrastructure, for instance where different protocols provide different transmission capabilities. However, the scenario primarily envisaged is that the network for carrying voice signals comprises a voice

10 network of the telecommunications type, providing substantially real-time voice communication over channels for which a connection (established for instance by maintaining route and/or bandwidth but more usually route) is guaranteed (within practical limits) for the course of a communications session.

15 Preferably, the network for carrying voice signals is provided by a public network such as the Public Switched Telecommunications Network (PSTN) in the United Kingdom. This gives very broad accessibility to services provided by embodiments of the present invention. The network for carrying data meanwhile may be provided by any appropriate network, such as a Local or Wide Area Network, or by

20 a complex of interconnected networks such as that known as the Internet.

25 Preferably, the means for receiving service requests is able to receive service requests at either one of the connection channel and data network interfaces. This allows users of either type of network to access services provided by use of embodiments of the present invention.

30 Preferably, any of a plurality of different applications may be run, depending on the service requested, and the services may be provided over one or both of the voice-supporting and the data networks. Thus embodiments of the invention can provide a powerful and flexible bridge for service requests coming in from either network and being provisioned over whichever network is relevant. For instance, Internet service requests can be provisioned over the PSTN and vice versa. Further the nature of the service may be simple or complex, being determined by the individual applications and the resource availability.

For instance, in an embodiment of the present invention, if a service request is received via the connection channel interface, it will usually comprise a telephone number and it will be received on one of the connection channels. Preferably, one or more digits of the telephone number represent the service being requested. An applications platform according to an embodiment of the present invention will then refer to the interpretation means to identify the relevant service and thus the application(s) required to support the service.

10 Such an embodiment has the advantage that a service request incoming from the voice carrying network can come in on any of several connection channels. This is not the case where it is the incoming channel itself which identifies the service required. If the service request can come in on any of several connection channels, then better use of the connection channels can be made.

15 Preferred embodiments of the present invention are provided with resource allocation means. The applications will then need resources such as connection channels, processing capacity and algorithms in order to support services. The resource allocation means can identify the application required to support a service, for instance by reference to interpretation means as described above, and subsequently allocate resources for running the relevant application.

In these preferred embodiments, it is advantageous if the resource allocation means can make use of "Direct Dialled Digits". That is, the user dials a telephone number which comprises the number for the applications platform, plus one or more extra digits identifying (with the help of the interpretation means) the required application. This is in the manner of dialling in to an extension number on a private branch exchange. The incoming service request in this case can access the applications platform by any of the plurality of connection channels. The one or two extra digits dialled by the user comprise a "Direct Dialled Digits" field which the resource allocation means can process to identify a relevant application by reference to the interpretation means.

By providing the resource allocation means with such data processing capability, and by including connection channels as allocatable resources, it is possible to allocate in real time any of the plurality of connection channels to an application necessary to provide a service, rather than having to pre-allocate channels. This 5 flexibility in resource allocation is extremely advantageous, particularly as the number and types of services to be provided, and therefore of the different applications which need to be launched, increase.

Preferably, the resource allocation means can operate in at least one other mode in 10 allocating channels to applications. For instance, a user may make a request for a service by dialling the telephone number. The telephone number may be specific to a particular incoming channel. The interpretation means, in order to support this mode, stores the relationship between incoming channels and applications required to supply the requested service. The resource allocation means then can refer to 15 the interpretation means to identify which application is relevant to that particular incoming channel. This provides interpretation of the request for service into a requirement for the relevant application by simply looking at which channel the request came in on.

20 Clearly, it is preferable that the applications platform is provided with means for receiving service requests at the interface for providing access to the data network.

The interface to the data-carrying network may advantageously be adapted to 25 receive electronic mail messages. For instance, it may be provided by a gateway such as an SMTP gateway. In this way, users may request services by means of sending electronic mail messages to the applications platform.

In preferred embodiments of the invention, service requests incoming from the 30 data-carrying network comprise electronic mail messages, each message comprising an identifier for a respective service and data for use in running the service, and there is further provided a data store for storing said electronic mail messages, at least one of the computer-based applications having access to the data store and having means to select a mail message carrying an identifier for a

service the application provides and to use the data from the selected mail message to provide the service.

For instance, a mail message may incorporate an identifier for a paging service.

5 The data it will carry will then usually comprise an identifier for a party to be paged. An application for use in providing the paging service will select the mail message because it carries the paging service identifier. It will then read the identifier for the party to be paged. The application will then act on the data to run the paging service to the relevant party.

10

Preferably, mail messages sharing a common service identifier may have more than one structure and the applications which select and act on those mail messages will have means to recognise the relevant structure and to respond by providing the relevant service appropriately.

15

For instance, a first mail message structure may be as described above, incorporating (at least) a specified field for the service identifier and a specified field for an identifier for a called party. The relevant application will provide the service identified, to the called party. A second mail message structure may then 20 have a further specified field, for instance for "unstructured data", the application being triggered to access the further specified field by a predetermined characteristic of the mail message. For instance, the pre-determined characteristic may be that the specified field for the identifier for the called party is empty.

Preferably the application has access to parsing means to parse unstructured data 25 in the further specified field. This allows the further specified field to be used in a particularly flexible manner in the requesting of services.

30 Preferably, as mentioned above, the applications platform is provided with resource allocation means to allocate resources to an application for use in providing a service.

Such allocation may be made in advance in relation to the running of an application, such that the allocated resource is unavailable to any other application, during the running of that application. Preferably, however, the resource allocation

means is equipped to receive and respond to a request from an application for the allocation of one or more resources, during the running of the application.

This has the advantage that resources that rely on relatively high processing

5 capacity but which are only required for a relatively short proportion of the running time of an application, such as speech recognition algorithms, are not rendered unavailable to other applications throughout the running time of any single application.

10 As mentioned above, the provision of interpretation means to relate applications to service requests allows flexibility in the allocation of channels to applications. Preferably, the interpretation means, for instance an interpretation table, is updatable and, more preferably, it is updatable while the applications platform is in use. This can be done by storing the interpretation means as an editable file, for

15 instance in ASCII on hard disk, rather than in the memory of a PC. Although this may be expected to lead to a response bottleneck, it has been realised that the files concerned are relatively small and the bottleneck does not occur in practice. Also, currently available disk controllers are provided with local memory for simple "read" commands. Hence access to the interpretation means does not always

20 necessitate access to the disk itself.

"Channel" in this context means a communication path of a type used to carry speech signals for a user. Such signals are usually carried in channels where the route through a network is set up and maintained during the course of a

25 communication session between the user and another network address such as a telephone number. "Channel" may therefore mean for instance a single time slot in a 30 speech channel ISDN connection. However, it is conceivable that other forms of channel which are capable of carrying speech might be used, such as a connection in which bandwidth rather than a particular route in a network is

30 reserved for the duration of a session.

A speech platform referred to as the "Minor Applications Platform" (MAP) will now be described as a specific embodiment of the invention, by way of example only, with reference to the accompanying drawings in which:

Figure 1 shows a schematic block diagram of components of the MAP;

Figure 2 shows a block diagram of the contents of a speech card as shown in Figure 1;

5 Figure 3 shows a schematic block diagram of two MAPs according to Figure 1, having some shared platform;

Figure 4 shows a functional block diagram of a MAP management system for managing overall process flow in the MAP of Figure 1;

10 Figure 5 shows a schematic representation of an example of data provision at startup for the MAP shown in Figure 1;

Figure 6 shows a flow diagram of events occurring in the MAP of Figure 1, at launch of an application in response to a service request;

Figure 7 shows a schematic overview of service and resource queues in the MAP of Figure 1;

15 Figure 8 shows a cyclic buffer for use in providing traffic statistics for the MAP of Figure 1; and

Figures 9 and 10 show in block diagrams two alternative approaches to providing multiple MAPs in order to support increased connections to the PSTN.

20 **GLOSSARY OF ACRONYMS USED IN THIS SPECIFICATION**

API	Applications Programming Interface
ASCII	American Standard Code for Information Interchange
ASIWR	Automatic Speaker Independent Isolated Word Recogniser
25 BT	British Telecommunications plc (the present applicant)
CDHMM	Continuous Density Hidden Markov Model
CSS	Customer Support Systems
DASS2	Digital Automatic Signalling System No. 2 (used in UK with provision of primary rate ISDN)
30 DDI	Direct Dialled In
DLIC	Digital Line Interface Card
DPNSS	Digital Private Network Signalling System (generally used between PABXs, has core capabilities and manufacturer - specific extensions).
DSP	Digital Signal Processor

DTMF	Dual Tone Multi-Frequency
DX CPU	DX Central Processing Unit
EMC	Electromagnetic Compatibility
GUI	Graphical User Interface
5 IDE	
ISA	Industry Standard Architecture
ISDN	Integrated Services Digital Network
LAN	Local Area Network
MAP	Minor Applications Platform
10 MVIP	Multi-Vendor Interface Protocol
NIC	Network Interface Card
NT	Trade Mark
ODBC	Open Database Connectivity
PABX	Private Automatic Branch Exchange
15 PC	Personal Computer
PCM	Pulse Code Modulation
PIN	Personal Identification Number
PSTN	Public Switched Telecommunications Network
SC	Speech Card
20 SIC	Service Identity Code
SMS	Short Messaging Service
SMTP	Simple Mail Transfer Protocol
STAP	Speech Technology Applications Processor
T120	An ITU standard
25 TCP/IP	Transmission Control Protocol/Internet Protocol
TLI	Transfer Line Identity
TTS(M/F)	Text to Speech (Male/Female)
VM	Voice Messaging
WWW	World Wide Web

30

Referring to Figure 1, the MAP 150 sits between a data network 155, such as the Internet, and a telecommunications network 160 such as the PSTN. It comprises primarily one or more speech cards 110 and one or more digital line interface cards 105, managed by applications running on a PC 120. A network interface card

(NIC) 130 provides connection to the data network 155, for instance by means of TCP/IP, and the DLICS 105 provide connection to an exchange 250, in this case a PABX, via an ISDN link 210.

- 5 The speech cards 110 carry a set of algorithms for detecting, recognising, recording and generating speech signals and tones. The speech cards 110 can also deal with modem signals. The PC 120 also carries some of the more sophisticated algorithms, such as the speech recognition and text to speech conversion algorithms. These are the resources which allow the MAP to recognise
- 10 and convert between different signal media, such as between spoken messages, recorded voice messages, electronic mail messages and the like. Additionally, there are one or more FAX cards 165 for providing the capability of sending and receiving by facsimile.
- 15 The network interface card(s) 130 may of course support different communication protocols. As shown, the network interface card 130 supports TCP/IP and provides an interface to the Internet 155. This can give access to services provided by the MAP to any user having a connection to the Internet 155.
- 20 The digital line interface cards 105, the speech cards 110 and the FAX cards 165 are connected to each other by a (known) multi-vendor interface protocol link 100 and to one or more personal computers 120 by ISA links 115.

- General control of operation of the MAP 150 is provided by a computer application
- 25 running on the PC 120. In particular, this application controls flow of processing in the MAP 150. Actual process execution to provide services is primarily (but not exclusively) carried out by processors of the speech cards 110.

- Referring to Figure 2, each speech card 110 carries seven speech processors. The
- 30 first of these, a Motorola 68,000 speech processor 205, provides processing power and gives overall control of what is happening on the speech card 110. The other six speech processors are Motorola 56,000 speech processors 200 which provide DSP capability. The 56,000 processors 200 provide some of the resources, particularly algorithms, used in running services on the MAP 150.

These algorithms provide functions such as MF4 detection, playing out messages at 8 kbits/sec, recording messages at 64 or 8kbits/sec, and speech recognition.

Overall control of the MAP 150 is provided by the "MAP Manager". This is

5 software which generally will reside on the PC 120. The MAP Manager has knowledge of the resources available, in terms of ISDN channels and the processors 200, 205, and it determines where resources will be made available to meet demand by applications. In particular, the MAP Manager has control over time switches 215 in the processors of the speech card 110 and allocates streams

10 for processing by means of the MVIP 100.

Referring to Figure 3, the basic architecture of a MAP 150 can be multiplied up so that there are at least two PCs 120, connected to their respective MAP speech cards 110. The DLICs 105 may share a common platform, as shown. Each PC

15 120 is connected via the ISA bus 115 to a network interface card 130 and thus to the Internet 155 via a network link 125. Further processing power can be provided by a UNIX processor 300, also connected to the ISA bus 115 and hard disk storage 305 is provided. Further memory 310 is associated with the PCs 120. In this arrangement, one MAP 150 can communicate with, and exploit

20 resources of, another MAP 150.

A MAP will now be described in more detail.

OUTLINE SYSTEM DESCRIPTION

25 Referring to Figures 1 to 3, the Minor Applications Platform (MAP) 150 described below is a PC-based speech platform capable of providing multiple speech applications on multiple telephone lines. The heart of the system is the speechcard SC2 110 which provides speech and signal processing based on algorithms. The

30 MAP 150 supports multiple digital links to public and private telephone exchanges using its Digital Line Interface Card (DLIC) 105. The MAP 150 supports resource sharing and is dynamically reconfigurable to run different applications on the same hardware. Internally the DLIC 105 and SC2 110 cards are interconnected using an industry standard speech bus 100 called Multi Vendor Interface Protocol (MVIP)

and this capability may be used to build MAP systems using other commercially available cards. MAP systems may also be built in shelves with passive backplanes and use plug-in processor cards. Such backplanes may be split in which case MAP hardware and software resources may be shared between 5 processors.

The MAP system has a single manager and can provide traffic statistics. Its operating system is UNIX. It has an archival facility which has direct links to DOS, WINDOWS and Netware.

10

THE NETWORK SERVER

The MAP system 150 provides a 30 channel digital connection 210 to a public or private telephone switch 250, such as a PABX, using a variety of protocols. The 15 MAP 150 allows connection to be made to a public exchange using signalling systems such as the DASS2 or A1B1 signalling systems, or to a private exchange using the DPNSS protocol. For the embodiment described below, the MAP 150 is connected to a private PABX using DPNSS which allows service provision based on TLI (and possibly SIC) codes. These codes have analogues in the PSTN where 20 the known C7 signalling system is used, and thus services provided on the MAP as described herein can be implemented over the PSTN.

In this implementation 3 PCs 120, 300 are used to provide the MAP system. One PC 120A provides system management and connection to the PABX 250, whilst 25 the other two PCs 120B, 300 provide advanced speech recognition and text to speech conversion services. A fourth PC (not shown) provides connectivity to the Internet 155 through an SMTP gateway, thereby enabling messages to be sent to and received from the LAN.

30 MAP Manager

The MAP manager provides control software for starting and terminating applications and co-ordinating their requests for services on line cards 105 and speechcards 110. The MAP manager resides on one of the PCs 120A and

provides a screen-based user interface. From the manager screen the MAP system can be started, parked or shutdown and the state of resources examined. The manager uses a small number of tables which may be examined to control allocation of resources and application selection.

5

A MAP "Traffic Statistics" screen can be invoked from the manager and provides basic information on platform use by applications, notably number of calls/application, average call holding time and busy hour.

- 10 Referring to Figure 4, the MAP Manager 400 comprises control software 440 and a set of agents 405, 410, 415, 420, 425. Each agent is generally allocated to a functional entity, which may be hardware and/or software, and provides a control link thereto. Hence there is a DLIC agent 405, a speech card agent 410, a fax agent 415, a text to speech agent 420 and a speech technology applications processor agent 425. The DLIC, speech card and fax agents 405, 410, 415 are provided with Applications Programming Interfaces (APIs) 430 and with plural drivers 435 for driving the relevant pieces of hardware (cards). Each API is provided by the suppliers of the relevant hardware cards. The MAP Manager control software 440 provides overall control of the agents and acts as a
- 15 sophisticated message switch in that substantially all messages go via the MAP Manager control software 440.
- 20

The drivers 425 receive incoming data and control outgoing communications for example by controlling interrupts on the relevant bus.

25

- 25 The MAP Manager 400 provides overall resource allocation. To do this, at start up, it constructs a world view. This is checked and updated at intervals. It obtains the world view through its resource agents. For instance, referring to Figure 5 which shows a representation of a small part of the MAP Manager's world view, the speech card agent 410 may have reported that two digital signal processors 500, 505 are loaded respectively with a multi-frequency detection algorithm 510 and an 8K coding scheme 515. These two digital speech processors 500, 505 will therefore deal with signals on the MVIP 100 which are incoming and outgoing respectively.
- 30

MAP Processor

- 5 MAP can be potentially hosted on most standard PCs with a 33 Mhz 486 PC or better and having a minimum of 8 Mbyte of memory. The standard processor used within the BT network is the Compaq 5133, a Pentium based machine having 5 full length ISA slots for MAP and other hardware cards.
- 10 As an alternative, a processor card manufactured by Aculab plc can be used. This is intended for use in a shelf unit where the need is to provide a multi-processor MAP system in a single enclosure with minimal non-MAP components. Outline features include
 - 15 80486 DX CPU operating at 25,33 or 50 Mhz
 - Internally clock doubled parts may be fitted
 - Pentium ready
 - 256 kByte cache
 - ISA Bus
 - 20 On-card Ethernet LAN interface
 - Serial and Parallel Ports
 - IDE Hard Disk and Floppy Disk Interface
 - Keyboard Interface

25 MAP Algorithms

- Speech algorithms for use on the speech cards 110 and processors 120, 300 include:

 - 30 64kbit/s speech I/O
 - 8 kbit/s speech I/O
 - Silence detection
 - DTMF Detection
 - DTMF Generation

Network Tone Detection

ASIIWR CDHMM Isolated Word Recognition Yes/No, Digits, alphas

ASIIWR Rapid Vocabulary Generation

STAP CDHMM Connected Recognition

5 STAP Rapid Vocabulary Generation

Fast/Slow Speech

Pulse Recognition

Noise Interrupt

Text to Speech

10 V21/V23 Modem

Audioconferencing

MAP Speech Card

15 The Speechcard is manufactured by Aculab plc and provides a UNIX-based environment for algorithm execution, with a VX Works real-time kernel. It contains six Motorola 56001A DSPs clocked at 27MHz for digital signalling processing each of which connects to the MVIP bus. The on-board 40 MHz Motorola 68000 coordinates communication between the DSPs and the MAP host processor over 20 the ISA bus. It is also used in some instances for further algorithmic processing.

The card is EMC approved and also provides a single analogue port for local speech recording/playback facilities.

25 MAP Operating System

The MAP operating system is System 5 release 4.4 UNIX. This is commercially available as Consensys, or Unixware V1.0. The MAP system could of course be ported to other systems such as Unixware V2.0 or LINUX.

30

The MAP system must also interwork with the speechcard host operating system, VxWORKS, a UNIX system with a real-time kernel. One copy is required per speechcard.

MAP Shelf Unit

5 This is manufactured by Aculab plc and contains an 18 slot passive ISA backplane organised into four processor sections as 5 + 5 + 5 + 3. Processor sections may be interconnected to make larger processing units. It is primarily used to contain multi-processor MAP systems and archival systems and contains

10 • Mains Power Supply
• Good Ventilation with fan failure monitors
• Internal provision for 4 Winchester Disks
• Floppy Disk Access
• Twin DAT Drive Access

15

MAP Digital Line Interface Card

Primary Rate ISDN connection is provided using Aculab's "E1" card. There are 30 and 60 channel versions available on an ISA PC bus. This E1 card supports

20 different types of speech bus 100, including MVIP. It is EMC approved and is widely used and approved for host-independent connection in many countries worldwide.

In the MAP system, as discussed above with reference to Figure 4, DLICs 105 are

25 controlled by a DLIC Agent 405. Under the control of the Manager 440, speech from a timeslot on an ISDN link 210 is switched by the DLIC 105 onto an MVIP timeslot and stream and, using a similar switch on the SC2 110, the manager directs the speech to a given DSP 200.

30 Referring to Figure 10, DLICs 105 in MAP systems 150 can also be "daisy chained", allowing one MAP system to pass calls to another, only the first MAP system 150A being connected to the host exchange. This avoids a situation as shown in Figure 9 where, in order to get sufficient MAP service capacity, a customer is supporting direct ISDN connections to three different MAPs 150. This

might otherwise be the result where a customer has an ISDN connection for 30 channels on one link but one MAP system 150 cannot provide all the required resource. To "daisy chain" MAPs 150, it is simply necessary to substitute a sixty channel DLIC 105 for the thirty channel DLIC 105 used normally in a MAP, in the

5 first and second MAPs 150A, 150B, and to use the second and third MAPs 150B, 150C for overflow connections via the DLIC agents 405 of the first and second MAPs 150A, 150B.

TRAFFIC MANAGEMENT

10

Referring to Figure 8, a particularly useful feature of the MAP 150 is that traffic statistics can be viewed substantially in real time on the MAP manager screen. The traffic management system 825 analyses event data for the MAP 150 by monitoring the MAP message stream in a cyclic buffer 800. It has visibility of the

15 interpretation tables available to the MAP 150 and can therefore identify applications from the messages.

In more detail, the traffic management system 825 sees all messages coming into the MAP by having a read access point 820 to a slot 805 in the cyclic buffer 800.

20 Every new message 810, or event, is stored in the next empty slot 805 of the buffer 800 and moves round as the MAP manager picks messages out of the buffer 800 at an output point 815.

The traffic management system 825 has a statistics stream processor 830 which

25 looks at the events 810 to detect traffic-related data such as incoming call flags or call terminations. It also has means 835 to access the interpretation tables and an accumulator 840 so that the traffic-related data can be organised against applications.

30 It is clearly necessary to choose a suitable cyclic buffer 800 arrangement, particularly in terms of the rate at which the MAP manager removes messages and the rate at which the traffic management system 825 reads them, which will not cause a bottleneck in the MAP 150. The processing load should be monitored so that modifications can be made as necessary.

The accumulator 840 allows the traffic management system 825 to look across multiple channels carrying calls and to provide a virtually instant traffic assessment. It can also provide analysed data in terms of for instance busy hour traffic, total number of calls, average holding time and customer-related statistics such as charges incurred by users.

THE CLIENT TERMINAL

- 10 The network terminal as described above provides the capability of the MAP to provide services which bridge between a switched voice-carrying telecommunications network such as the PSTN and a primarily data and message-carrying network such as the Internet. It can be accessed either by a telephone call via the voice-carrying network or by an access point such as a client terminal attached to the primarily data and message-carrying network. A suitable client terminal is described below.

The client software sits on a user's terminal 315 which could access the MAP in any of several different ways. For instance, it could have a socket connection to the MAP (not shown), a LAN connection (not shown), or it could be connected to the Internet 155. It is GUI-based and consists of up to two main components. It is intended for use on PCs running WINDOWS '95 or an NT workstation. In addition it may be possible for the software to run on WINDOWS 3.1 or with Windows for Workgroups V3.11 where the WIN32S extensions from Microsoft will be required. One software component is preferably a Microsoft "wizard", concerned with creating service requests via the Internet. The other software component is concerned with responding appropriately to inbound requests which may be for instance incoming information or connection requests.

30 Service Request Wizard

This facility, in the form of a Microsoft 'Wizard', guides the user through the necessary steps required to establish a service request. At each stage, the user is able to backtrack and change selected options. The number of steps required

before the service request is specified will, of course, depend on which service is requested. In using the services, the user will select the person with whom he wishes to communicate. The various numbers which are used to effect the communication can then be derived from a directory whose content is in part

5 established at registration time (see "REGISTRATION AND DIRECTORY SERVICES" below). For example, in establishing an audio conference the telephone numbers from a directory can be used, whereas had 'paging' been requested individual paging numbers would have been entered by each registrant.

10 A disadvantage of this approach is that identical messages for dispatch by different mechanisms have to be re-entered. An alternative approach is to first select the recipients and then, for each recipient, to select a method for receipt whilst taking account of known options for each person.

15 It should be noted that a single service request does not necessarily imply use of a single backbone service. For example, to establish an audio-graphics conference requests may be made of both audio conferencing services and T120 services.

Connection Request

20

This software, normally iconised and possibly invisible, will respond to incoming messages from the MAP system. Such inbound messages might occur for two purposes:

25 1. Information. For example, it may not be possible to provide a requested service because, perhaps, a telephone line is engaged or unanswered after a predetermined number of attempts, in which case the requester needs to be notified of the circumstance.

30 2. Connection Requests. These will typically arise as a result of an Internet service request by another party; for example, it may be an invitation to participate in an (Internet) video or T120 conference. The action should be to give the user the opportunity to accept/reject the request and, if accepted,

to launch the (requested) service on the user's terminal. In the above example, this would be either 'CUCme' or 'Farsight'.

MAP SERVICE PROVISION

5

Referring to Figures 3 and 7, requests for services using the MAP may be received either on one of the ISDN links 210 or, via a network interface card 130 and an SMTP gateway process 725, from the Internet 155.

- 10 To respond to a request for service, the MAP 150 will need to run a computer-based application 730,735 to manage the service, and it will need to allocate the resources necessary to run the service. Usually, these resources will include at least one ISDN time slot 720. Service requests coming in on an ISDN link 210 will necessarily already have an ISDN time slot 720 and the MAP allocates these directly to the relevant applications 730, 735. Service requests coming in from the Internet on the other hand are placed in a Generic Service Queue 710 and the various applications 730, 735 which are relevant to Internet-requested services read out their own service requests from the Generic Service Queue 710.
- 15
- 20
- 25

attempt should be replaced by fax notification.

Applications 730,735 on MAP speech channels 720 can in general be caused to run in two ways.

- 30 1. Immediately on system startup. This mode is used, for example, as the basis for provision of Internet requested speech services. If a speech service request is received by Internet, there is no speech channel 720 directly involved. One needs to be provided as quickly as possible. A minimum of channels 720 are therefore permanently dedicated to each service, though an

application can acquire further lines from a free pool. Such applications 730,735 are normally outgoing only e.g. notification and voicemail delivery services, or booked services such as conferencing.

5 2. On receipt of incoming call. In this case, there is necessarily a speech channel 720 associated with the service request. The MAP system allows the application to be selected by channel of entry or DDI number, both in conjunction with the SIC. In the present embodiment, the DDI method has been adopted as this allows line plant to be used according to demand with 10 consequent higher line utilisation. This facility would be used for services such as telephone access to e-mail.

As mentioned above, more sophisticated services may be conveniently built up on more basic services - for example, should it be desirable to send a FAX as part of 15 one service, the document for transmission can be sent to a FAX service for transmission to the intended recipient. Such delivery services are non-real time and by their nature do not call upon other services. Requests for such services are queued, with the relevant MAP applications taking requests from the generic service queue 710.

20 The Generic Service Queue 710 can be divided into application types so that there are multiple service queues 710. A number of service channels can be pre-allocated to each service queue 710 and the number of service channels for a particular service queue can be determined from observed traffic levels.

25 An Internet address can be associated with each generic service queue. Requests can then be added to the queue by sending messages to this address. These messages need to contain all the information required for the MAP application to subsequently enact the request, and are thus best computer generated. Client 30 software to originate such messages is described above, and is the basis on which 'point and click' telephony and messaging services can be provided. However, it should be noted the MAP applications can themselves request such backbone services. For example, if a user of the message collection service wants to receive

a copy of an e-mail message by FAX this can be achieved by the voicemail application making a FAX request. This is the basis of the collection services.

MAP Resource Allocation

5

This is further described below. However, resource allocation in general is as follows:

- *Line Allocation* is specified in configuration files which may be dynamically updated while the system is running. Lines may be classed as reserved for outgoing (free pool), fixed for incoming, or available for either. Lines may be allocated either statically or dynamically.
- *Application Allocation* is either by point of entry (ordered) or by DDI digits for service requests incoming by ISDN channel. For incoming SMTP messages, it is simply done by interpreting the service requested and directing the request to an "IMMEDIATE" application. Applications are normally run by the manager but, like hardware resources, may be distributed between MAP systems sharing a common MVIP.
- *Algorithm Allocation* may be either static, dynamic or a combination of both. There can be only one algorithm per DSP, but (depending on the algorithm) one DSP may be able to serve more than one channel. There is no fixed algorithm/DSP association and this will typically change according to the demands of applications.
- With *static* resource allocation, all resources are allocated at the time of call arrival and will be always available even if unused. With the *dynamic* scheme, resources are only allocated when needed which enables more channels to be supported for given level of DSP resource - there is however a probability of demanded resources being unavailable which must be suitably catered for within the controlling speech application. This situation can be handled in MAP by the application requesting resources in advance and manually releasing them when no longer required.

INCOMING PSTN CALL PROCESS

Referring to Figures 3 and 4, a user initiates a dialled connection to the MAP 150, using a telephone 320 connected to the PSTN 160. The call will be received at a

5 DLIC 105 and a driver 435 for the relevant DLIC agent 405 will receive an event on a timeslot "X". The DLIC agent 405 has its own processing capability and will instruct the driver to acknowledge the event on timeslot "X". The DLIC agent 405 will have also instructed the relevant DLIC 105 as to what protocol is being used. The DLIC agent 405 then tells the MAP Manager control software 440 that there

10 was an incoming event.

The MAP Manager control software 440 recognises that it will have to raise a computer-based application to provide the requested service. To do that it needs to obtain the DDI digits which will identify the service requested by the user. The

15 MAP Manager 400 will therefore tell the DLIC agent 405 that it needs the DDI digits.

Once the MAP Manager 400 has the DDI digits, it will access an interpretation table stored on disc 305. In practice, the disk 305 will store any frequently

20 accessed data in local memory, for read access only. (Nowadays this is a standard feature of many disk storage systems. It is controlled by a disc controller provided by the disc supplier.) The DDI digits, once interpreted by use of the table, will identify to the MAP Manager the service required. The MAP Manager 400 needs to translate this into the application(s) needed to provide the service, and thus into

25 resources, including algorithms, required. It therefore looks up another interpretation table, in the same manner as before, in order to find out what resources and algorithms are going to be needed to run the applications necessary to provide the service. (In practice, one interpretation table may provide both service and resource data.)

30

Knowing the resources, including algorithms, it will need, the MAP Manager will now look at its own world view, constructed on start-up, to find a functional and available digital signal processor 200 and spare ISDN channels 210, with which to run the relevant application. Having found the relevant resources, the MAP

Manager 400 books the resources by entering a flag in the internal table which is the world view and sends messages back to the relevant agents 405, 410 to allocate switch settings 215 and to launch the relevant application.

- 5 Referring to Figure 6, the first task of the application is usually to answer the incoming call. The application 600, using the relevant API 605, now tells the MAP Manager that the call should be answered. The MAP Manager 400 tells the DLIC agent 405 to answer the call and the DLIC agent 405 passes the message on to the application via the MAP Manager 400. The MAP Manager 400 now reports 10 back to the application 600 that the call has been answered or that there has been a failure. At this point, the user who has initiated the call hears the ring tone leave the line.

The application next instructs the MAP Manager 400 to play a file. The MAP 15 Manager 400 does this by means of the speech card 110 and by passing a request to the SC2 agent 435. Once notified that the file has been played, the MAP Manager 400 sends notification back to the application 600.

The above describes the first steps in the running of an application 600 to provide 20 a service using the MAP 150. The application 600 will continue to control provision of the service until it has concluded. At this point, the MAP Manager 400 will send an acknowledgement back to the DLIC agent 405 and the DLIC agent 405 will clear the line. This means the switches 215 will be opened and the resources will be "torn down", for instance by marking the digital speech 25 processors 200 as available again in the internal table of the MAP Manager 400.

It will be recognised that the DLIC 105 carries relatively little data. There is significantly more data on the speech cards 110. The 68,000 processor on a speech card 205 runs the VX Works UNIX kernel 170 in real time, by means of 30 which it can run VX Works. If the speech card 110 receives a play request from the MAP Manager 400, it is VX Works which handles it. That is, it runs a file off disk without involving the main processor. (This is a known use of UNIX).

It is possible to do the messaging between the MAP Manager 400 and its agents 405, 410, 415, 420, 425 via common memory or via a network interface card 130. VX Works will run as though it were stand-alone, using a network interface card 130.

5

It will be recognised that although the speech cards 110 do a significant amount of processing, processing can also be distributed elsewhere. A constraint is that although applications can sit in many different places, all speech signals should be carried via the MVIP bus 100. Consequently, speech databases should sit next to 10 their speech cards 110 to avoid carrying speech signals on the Internet, or other relatively unsuitable highway.

INCOMING SMTP MESSAGE PROCESS

15 Referring to Figure 7, service request messages enter the system through the SMTP gateway 725, the message request being physically transferred to the system and (usually) placed on the relevant generic service queue 705, 710. This gateway 725 relates Internet addresses to generic service queues 705, 710. The addresses of the available generic service queues 705, 710 are of the form

20

<user>@<service>.map.cartoon.bt.co.uk

where <service> is the name of the generic service queue. It should be noted that the <user> field may not be used or may have different interpretations in 25 different services. Messages to non-existent services are discarded. Messages to services which require immediate delivery will be sent directly to the relevant service channel queue 730, 735.

Client Services

30

In order to request core services, the Internet needs to originate a suitably structured message. This can be created by hand and dispatched using any available mail tool, but such an approach can present difficulties.

- The message may be wrongly formatted or incorrectly addressed, in which case the request may be at best ignored.

5 • The message may take a significant period to propagate from the customers mail system to the MAP's SMTP gateway 725, and thus make the service potentially untimely.

- The system would be inconvenient to use.

10 A better approach is to create a simple program for execution on the client which creates and properly addresses the request. Client SMTP software can be embedded in the application enabling fast and direct client - MAP SMTP communication and thereby provision of service. This program can take the form of a 'Wizard' taking the user through the necessary steps to make his request -

15 this may only be a few clicks of a mouse, and the user need know nothing of the mechanism by which the requested service is provided.

..... This functionality can be provided on a Worldwide Web (WWW) server, in which case the client can be any WWW browser.

20 An important part of such client software is the ability to be able to select intended recipients for the communication. The software can conveniently provide ODBC links to directories used within a customer's private network or Intranet. These directories may be either local to a user's hard disk or LAN, or potentially on

25 remote computer systems. Use of these directories will enable lookup of telephone numbers etc and thus, in conjunction with the above program, enable 'point and click' communications to be implemented.

..... Some backbone services potentially require an element of communication with the

30 requesting user. This might be to provide a booking reference number, to advise of problems in being able to contact some parties, or perhaps to advise of an inbound communication. Thus the MAP SMTP gateway 725 must be able to create/forward outbound messages from the MAP platform 150 and some client software is also desirable to receive and act upon such messages. This client

software would normally be iconised until activated by message arrival when it should interrupt the user to request what action should then be taken. This requires extensions to the MAP SMTP gateway 725 and client software to act on such messages. (Client software is already available which could be easily

5 modified to accept such messages.)

In the normal course of operation, the MAP creates files of raw data on service instances it handles. These are each identified by the "service_instance_identify" field (See "Protocol for Services on MAP" below) and stored in a raw call record

10 data area (an FTP directory) of a database. It could be on the same computing platform as the MAP 150, or it could be elsewhere.

The MAP is also able to output such files via an FTP link if necessary for processing remotely, for instance by a user's remote application.

15

So far as launching applications and delivering resources for services is concerned, the process is the same where a service request has been received at the SMTP gateway 725 as described above for the "Incoming PSTN Call Process". That is, the MAP manager 440 will use its various agents as described above with

20 reference to Figure 4.

PROBLEM SCENARIOS

User Rings Off

25

The above describes service provision which runs to its normal end. It may be that the user rings off while a file is still being played. In this event, the DLIC agent 405 will notify the MAP Manager 400 which in turn notifies the speech card agent 410 that the file should be stopped. The speech card agent 410 will tell the

30 relevant speech card 110 to stop the file and report back to the MAP Manager 400 which in turn sends an acknowledgement back to the DLIC agent 405, and notifies the application that the line has cleared, once any house-keeping activities are completed, such as billing. Once complete, the application terminates which results in tearing down the resources, as before.

Wrongly Distributed Resource

There are other problems which may arise. For instance the MAP Manager 400

5 may find that it does not have sufficient resource, such as "play". In these circumstances, the MAP Manager 400 will seek to change the available resources and therefore to change its world view. It may do this by deleting "recognition" from a processor, as long as it is not in use, and substituting "play" resource.

10 Resource Failure

In a different problem scenario, there may be functional difficulties, for instance with a digital signal processor 200. Say a speech card 110 was providing a play file. This is done in discreet amounts. The digital signal processor 200 should

15 respond in a given time (to VX Works) to say it has played a required section of a file. If a digital signal processor 200 has "died", VX Works will never receive an answer. In these circumstances, the speech card agent 410 will reload the code to the digital signal processor 200. It will, in practice, try this reload three times. If the speech card agent 410 has not received any response from the digital signal 20 processor 200 after three reloads, it will report the digital signal processor 200 to the MAP Manager 400 as being of doubtful integrity. The MAP Manager 400 will select another DSP 110 and load the relevant algorithm. The MAP Manager 400 will subsequently only use the digital signal processor 200 marked as of doubtful integrity in emergencies.

25

If then another digital signal processor 200 fails, the MAP Manager 400 has a potentially serious situation. It now issues a message to the control console (accessible from the MAP Manager screen), and issues a serious fault indicator on the serial line for that purpose. It may not be the digital signal processors 200

30 which are failing but VX Works for instance. The MAP Manager 400 will also take the relevant card or cards out of service and use spare resource. At this point the MAP Manager 400 reboots the suspect card. This forces the card to go through a "power-on self test". The card will then be proven either serviceable, in which case the MAP Manager will put it back into use, or failed.

Certain applications require large amounts of resource. For instance, speech recognition would create a bottleneck for the digital signal processors 200, 205. Referring to Figure 4, the text to speech (TTS) agent 420 accesses resources in a 5 dedicated processor via a network interface card 130. The TTS agent 420 will send a sentence to the dedicated processor. It receives back a 64K speech file which the TTS agent 420 runs and plays through a digital signal processor 200.

Running the digital line interface card 105 is relatively easy in view of the lower 10 data quantities. More of the management effort tends to lie with the speech cards 110. The processor on the MAP host, running UNIX, is usually under-utilised. This processor can be used to host text to speech or STAP algorithms. It should be noted, therefore, that the TTS agent 420 is not controlling hardware in the manner of the other agents 405, 410, 415, 425 but is controlling processor 15 resource. The TTS agent 420 therefore has to allocate resource in a different way. The speech card agent 410 measures resource in terms of channels on the digital signal processors 200 and how much 68K processors hire an application will take. For instance, for multi-frequency tone recognition, ten channels would need to be allocated and 2% of a UNIX controlling process. Looking at the TTS agent 420, it 20 is necessary to look at processing resource and it will be measured in megabytes and channels, for instance 32 megabytes and seven channels. The TTS agent 420 is directly connected to the speech card 110 by socket, this providing a real time connection. The MAP Manager will therefore instruct the speech card agent 410 that it needs 64K play capacity and a socket interface. TCP/IP runs on the 25 Ethernet at about 2 megabits per sec which is the equivalent of about 27 channels. Hence the socket arrangement avoids a bottleneck at the PC bus point which would only be able to provide about 12 - 18 channels off disk.

The STAP agent 45 provides the inverse process to the TTS agent 420. Speech is 30 received through the speech card agent 410 which does some pre-processing. It then sends speech to the STAP agent 425. The STAP agent 425 returns a recognition result and the speech card agent 410 notifies the MAP Manager 400 accordingly.

BACKBONE SERVICES

5

These services are important in enabling more sophisticated services to be provided, incorporating within them one or more backbone services. Four backbone services are described below. These can be addressed by client software. (Other "backbone" services may also be provided to enable facilities such as user registration and call accounting).

E-Mail Delivery Service

The queue for this service takes as its input a text message, and attempts to deliver this message by speech to one or more specified telephone numbers. The message may be delivered by the voice determined by the Internet address chosen according to the following table:

Internet Address	Voice
<user>@ttsm.map.cartoon.bt.co.uk	British Male
<user>@ttsf.map.cartoon.bt.co.uk	British Female

20 The service comes in two forms. If the user field is present in the address and is all digit, it is interpreted as the telephone number which should be dialled from the United Kingdom (UK) to contact the (single) receiving party, with the body of the message being taken as the unstructured whole of the communication. If the user field is absent, the message body is taken as being structured potentially to enable 25 distribution to many parties.

Three attempts should be made at delivery of a message to a given party and, if the message cannot be delivered, a notification message should be dispatched to the originator giving the reason for non-delivery. If the target telephone line is 30 answered, there is a simple service announcement followed by a query as to whether the message can be delivered. Once delivered, the customer can have the

31.

option to hear the message again or to receive a FAX copy. If delivery was turned down, a simple message should be sent to the originator detailing the non-delivery.

Paging and Short Messaging Service

5

The queue for this service takes as its input a text message, and attempts to deliver the message as text using the service determined by the Internet address chosen according to the following table:

Internet Address	Service
<user>@page.map.cartoon.bt.co.uk	National Radiopaging
<user>@sms.map.cartoon.bt.co.uk	GSM Short Messaging

10

Each service comes in two forms. If the user field is present in the address and is all digit, it is interpreted as the UK paging number to contact the (single) receiving party, with the body of the message being taken as the unstructured whole of the communication. If the user field is absent, the message body is taken as being 15 structured potentially so as to enable distribution to many parties. It should be noted that for the SMS service a 160 character limit is imposed, and a smaller limit may apply to the National Radiopaging facility.

FAX Service

20

The queue for this service takes as its input a text message, and attempts to deliver the message by FAX using the service determined by the Internet address as specified in the following table:

Internet Address	Service
<user>@fax.map.cartoon.bt.co.uk	FAX Messaging

25

Three attempts should be made at delivery of message to a given party and, if the message cannot be delivered, a message should be dispatched to the originator giving the reason for non-delivery.

Telephony and Audio Conferencing Service

The queues for these services take as their input a structured text message which specifies the telephone number of the originator and called party(ies), and

5 subsequently attempt to establish telephone calls to the people involved. The Internet addresses for the services are given in the following table:

Internet Address	Service
<user>@tel.map.cartoon.bt.co.uk	Telephone Call
<user>@conf.map.cartoon.bt.co.uk	Conference Call

Telephony could be treated as a conference call between two people and may be

10 implemented in this way, but in practice a separate queue is established as in this instance the connection can be made without recourse to use of the conferencing algorithm. In a conference call the message originator is considered to be the chairman. Three attempts are made to connect to each party and the participants notified by the system of progress in establishing the audio conference, with

15 already connected parties being able to converse.

REGISTRATION AND DIRECTORY SERVICES

Registration Services enable a user to initially lodge detail of his communications

20 methods, preferences and services required. For example, he might elect to take the PhonE-mail service and request notification about new messages at 10.00 pm by paging. In this case, he would declare his telephone number, paging/SMS numbers, any mobile phone numbers, TCP-IP address etc., and in this way a personal profile for the user would be established. This registration is achieved via

25 a WWW page, and following successful registration the user is issued an account number and a PIN. The account number/PIN would be used to access the service(s) either via WWW or telephone. On WWW the user would have the opportunity to revise his communications and service selections as his needs determine.

Information provided at registration is also made available in directories. These directories may be available on networked fileservers and accessed via trusted GUI front-end software. Typically they would be used to select people with whom the user wishes to communicate.

5

INTEGRATED SERVICES

Applications may be written for the MAP which provide integrated services by appropriately calling upon directory and backbone services to provide the required

10 overall service. Such services come in three main types:

1. Collection Services where the user connects to the system to receive any messages that may have been delivered.

15 2. Notification Services which notify the user of any messages waiting.

3. Delivery Services in which messages received are delivered to the recipient in an appropriate way.

20 Each of these will now be outlined in turn with reference to a network-based messaging system capable of accepting e-mail and telephone messages.

Collection Services

25 The customer may connect to the system either speculatively or because he has been notified of messages waiting. The customer will normally be invited to enter an account number and a PIN (issued at time of service registration) to identify himself to the system following which he is offered the various messages held on the system. The way in which he is able to select messages and have them

30 passed to him is a feature of the application providing the collection service, and internally to provide the various facilities will dispatch various messages onto the core service queues. For an outline description of one such service, PhonE-mail, see below.

Notification Services

The SMS, paging and TTS systems may be used to notify users that they have received messages - these may be either telephone messages, electronic mail

5 messages or potentially FAX messages. The notification message can advise the customer how many messages of which type he has received, the message itself being time/date stamped. By reference to the customer's service profile, notification of message availability would occur at a time of day determined at registration by the customer rather than at the time of message arrival since BT

10 may wish to make notification a subscription service.

It should be remembered that notification services only advise of message arrival. Storage of messages may make large demands on available disks, and for this reason it may be necessary on systems such as 'PhonE-mail' to limit the amount of

15 storage either by size or duration e.g. messages may only be held for 5 days. Thus, when a customer comes to collect his message(s) he may find fewer available than were reported by an earlier notification message.

Delivery Services

20

Delivery Services which could be offered come in two forms:

Immediate Delivery. In this case, on receipt of an inbound message and by reference to the customer's service profile a message received on one queue could

25 be re-directed to another queue. For example, in his profile a customer may have directed that all e-mail and voice messages would be delivered by female voice to a specified telephone number. In this instance all messages on the 'email' queue would be reformatted and placed on the TTSF queue and telephone messages placed on the 'VM' queue. Both new messages would be marked for immediate

30 delivery. It is assumed that if the customer has chosen immediate delivery there will be no notification.

Deferred Delivery. The delivery mechanism is the same as for immediate delivery above, except that the messages placed on the queue would be marked for

delivery at a time determined from the recipients profile. Delivery as an adjunct to notification is to be preferred as otherwise all pending messages will be scheduled for delivery at the same time.

5 PHONE-MAIL

The PhonE-mail service allows a customer to redirect his telephone and email in such a way that he can later collect both forms of messages, by ringing a particular telephone number, from any kind of telephone.

10

On service registration by WWW he will be given an account number and a PIN number and will be able to specify his communications profile, service requirements and messaging capabilities. He can later change these options by connecting to another WWW page showing his account details (following entry of

15 account number and PIN). Typical entries would include his normal telephone number and/or mobile telephone number and include his SMS number and pager number. If he selects notification he would enter a time of day or perhaps request immediate notification.

20 When using the service he would redirect his e-mail to an Internet address. This would take the form:

Internet Address	Service
<useraccount>@email.map.cartoon.bt.co.uk	Electronic Mail

For example, if at registration the customer was given number '9178' he should
25 redirect his mail to the PhonE-mail at address 9178@email.map.cartoon.bt.co.uk. As for other core services, this would be directed by the SMTP gateway onto queue 'email' but be processed by the associated MAP application. On this queue messages are unstructured, i.e. do not have the form outlined above, and will be parsed to identify relevant information concerning the originator, the most
30 important elements being his name, the subject of the message, and the body of the message itself. This text processing will also remove, as far as possible, all unwanted and irrelevant information from message headers often inserted by mail

systems. The text from each of these fields will then be converted from text to speech and the resulting text and speech files stored in the customer's account on disc. The speech duration of the message body should then be compared against a predetermined threshold which, if exceeded, would result in the text body being 5 passed through a summariser and the resulting text also being converted to speech. At this point the e-mail message is considered to be available for collection by telephone.

In a similar way, the customer may redirect his telephone to a number made 10 known to him at registration. Calls to his normal number will then be directed to the MAP platform which can, by extraction of the TLI, identify for which user the call was intended. Following a suitable greeting, the caller can then be asked to leave his name (and contact telephone number), the reason for his telephone call, and a message. These three aspects correspond to the name, subject and body of 15 an e-mail message. The speech can be stored along with that derived from e-mail messages, and the telephone message at this point may be considered to be available for collection.

Once a message is classed as available, the customer profile should be checked to 20 see whether he should be notified of message arrival, and if so how it should be achieved. If message notification is requested, it may be either on message arrival or at a particular time of day in which case a flag should be set to show messages pending for that customer. That flag would be interrogated by a regular scheduled process. The notification methods available can be achieved by one of the 25 backbone services - email delivery (male or female), paging or SMS. It would be possible to redirect email to FAX.

To collect messages, a customer would ring a designated telephone number. Following a simple greeting he would be asked to enter his account number and 30 PIN using the MF keypad. He would be given three opportunities to correctly enter the combination, and if unsuccessful perhaps invited to ring a helpdesk before the connection to the system is dropped by the platform. On entry to the system, he would be advised of the number (and type) of messages available.

The form of the ensuing dialogue could be of different types. For instance, the caller might be given the option of listening to the next message or reviewing his messages. Messages can be presented in order of arrival i.e., oldest first. Review options are by sender or by subject. Once a message is selected, the sender, the

5 subject and the message body will be spoken to the caller, though if the spoken length of the message body exceeds a threshold notionally set at 20 seconds the user will be queried as to whether he/she wishes to hear the entire message as it may be more appropriate to receive long messages by FAX. Once a message has been heard, the customer is given the option of being sent a FAX copy of the

10 message to a number he designates before choosing whether or not to delete the message. Once he has responded to these options, he can move on to listen to the next message, and so on until he has heard all the messages he wishes to hear. All announcements in the demonstrator are overridable by MF and the FAX/Delete options can be responded to by voice. Note that any messages not

15 deleted will be classed as new messages the next time the caller rings into the system.

Possible changes to the above include

20 • Presentation of messages in last-in first-out order i.e. most recent first.

• Modification of the dialogue to enable an experienced user to quickly recover all his messages

25 • Messages to be classed as either 'new' or 'saved'. New messages are messages which have not been heard; saved messages are messages which have been heard but not deleted

• All messages would be automatically deleted after a predetermined time e.g.

30 5 days

• Capability to send messages containing Microsoft 'Office' attachments by FAX

The application would run as a 'layered' service, using where possible the core services to provide the messaging required.

INVOICING AND CALL RECORDS

5

MAP applications can generate log files which describe actions taken by the platform for a particular call. For instance, to create the call record data area referred to under "Client Services" above, when the MAP is in the process of setting up a connection and transmitting a message to a recipient, it will output an

10 update status message to a raw call records data directory, that is the FTP directory, whenever there is a change in status for a service instance. It uses the "service_instance_identify" field (see "Protocol for Services on the MAP" below) to identify the relevant file.

15 Additional information may be held in these log files to provide information for billing purposes. For example, the 'servicename' and user account number can be included; this information could be included as part of the structured messages so that layered services such as 'PhonE-mail' which call upon basic services can pass the essential billing information forward in a way which is compatible with Internet
20 GUI originated messages. The log files once created would normally be dispatched to a specified mail address where they would be analysed to extract information for billing purposes. If this is to be handled on-platform the receiving process may take the form of a service queue. Potentially, billing information can be made available to the user via WWW.

25

PROTOCOL FOR SERVICES ON MAP

The following gives a specification of each of the variables in the MAP service messages so as to clarify the protocol between a client application and the MAP.

30

The message

Messages to the MAP will be addressed in the following form:

<user> @ <service> .map.cartoon.bt.co.uk

The user field

5 The <user> field may contain the telephone number of the person requesting the service. It will not necessarily be used in all services, but it is good practice to include it for all services.

The service field

10

The <service> field contains the name of the service being requested. Some examples are given below:

	ttsm	text to speech - male voice
15	ttsf	text to speech - female voice
	tel	telephony
	page	to send a page message (may need to split according to swatch, data pagers etc)
	sms	sends an SMS (Sort Message Service) to a GSM phone
20	fax	fax message
	conf	telephone conferencing

All services will exchange the same message format. Redundant fields will be assigned the string NULL. This makes it easier to code the client and server as well as making it easier to manage the interface.

In the specification of the messages 'service_instance' refers to a message or a conference.

30 Setup message

service_instance_identify =

- the identifier assigned to the service_instance by the MAP
(initially this will be 0000)

- start_date_time =**
 - the start date and time given as Mmddhhmm
- stop_date_time =**
 - the stop date and time given as Mmddhhmm
- 5 sender_telephone_number =**
 - the telephone number of the person setting up the service_instance including STD code
- sender_name =**
 - the real name of the person setting up the service_instance
- 10 sender_email_address =**
 - the email of the person setting up the service_instance
- sender_fax_number =**
 - the fax number of the person setting up the service_instance
- sender_pager_number =**
 - the pager number of the person setting up the service_instance
- 15 sender_SMS_number =**
 - the SMS number of the person setting up the service_instance (GSM phone number)
- number_of_receivers =**
 - the number of people receiving the service_instance (not including the sender)
- receivers_telephone_number =**
 - a list of the tel numbers of the receivers (separated by whitespace)
 - same order as receivers_telephone_number
- 25 receivers_email_address =**
 - a list of the emails of the receivers (separated by whitespace)
 - same order as receivers_telephone_number
- receivers_fax_number =**
 - a list of the fax numbers of the receivers (separated by whitespace)
 - same order as receivers_telephone_number
- 30 receivers_pager_number =**
 - a list of the pager numbers of the receivers (separated by whitespace)
 - same order as receivers_telephone_number
- receivers_SMS_number =**
 - a list of the SMS numbers of the receivers (separated by whitespace)
 - same order as receivers_telephone_number

41

- a list of the pager numbers of the receivers (separated by whitespace)
- same order as receivers_telephone_number

receivers_SMS_number =

- a list of the SMS numbers of the receivers (separated by whitespace)

5 - same order as receivers_telephone_number

confirmation_email_address =

- the address to which the confirmation that the service_instance has been successful should be sent

start_now =

10 - set to YES if the service_instance is to be started as soon as possible
- set to NO if the service_instance is to be started at the specified time

delete_service_instance =

- set to YES if the service_instance is to be deleted
- set to NO if the service_instance is not to be deleted

15 message_start =

- flags the beginning of the message body

message_stop =

- flags the end of the message body

20 Confirm/reject message

service_instance_identity =

start_date_time =

stop_date_time =

25 sender_telephone_number =

sender_name =

sender_email_address =

sender_fax_number =

sender_pager_number =

30 sender_SMS_number =

number_of_receivers =

receivers_telephone_number =

receivers_name =

receivers_email_address =

42

receivers_fax_number=
receivers_pager_number=
receivers_SMS_number=
confirmation_email_address=

5 start_now=

delete_service_instance=
message_start=
message_stop=
status=

10 - set to FAIL if the TTS service_instance has been rejected

- set to SUCCESS if the TTS service_instance has been accepted

15

20

25

30

CLAIMS

1. Apparatus for providing communications services in response to service requests, the apparatus comprising:

5

i) a connection channel interface, for connecting the apparatus to at least two connection channels of a network for carrying voice signals;

10 ii) a data network interface, for connecting the apparatus to a network for carrying data;

15 iii) a resource interface, for connecting the apparatus to resources for use in providing communications services, said resources comprising at least one speech-related resource;

15

iv) interpretation means for defining a relationship between service requests received at the connection channel interface and computer-based application programs for use in provision of the services specified in respective requests; and

20 v) initiating means arranged in operation to identify and initiate the running of one or more computer-based application programs in response to said service requests received at the connection channel interface, said program(s) being identified by reference to the interpretation means, which one or more computer-based application programs is/are adapted to call on at least one of said resources to 25 provide a communications service in response to a received service request,

such that the apparatus is able to respond to a received service request by identifying one or more computer-based application programs for use in provision of that service and initiating the running of the identified program(s).

30

2. Apparatus according to Claim 1, wherein service requests can be received at the connection channel interface and service requests can be received at the data network interface.

3. Apparatus according to Claim 2 wherein the initiating means is also adapted to initiate the running of one or more computer-based applications, in response to service requests received at the data network interface.
- 5 4. Apparatus according to any one of the preceding claims wherein the network for carrying voice signals and the network for carrying data are provided by the same network infrastructure.
- 10 5. Apparatus according to any one of the preceding claims wherein the network for carrying voice signals comprises a telecommunications network, providing voice channels in which one or more connections are established and maintained to support a communications session throughout its duration.
- 15 6. Apparatus according to Claim 5 wherein each of said connections is established over a fixed route for the duration of a session.
- 20 7. Apparatus according to any one of the preceding claims wherein the initiating means has access to a plurality of different computer-based application programs such that any one of a plurality of different communications services can be provided in response to a received service request.
- 25 8. Apparatus according to any one of the preceding claims which further comprises queuing means for queuing received service requests, said initiating means responding to received service requests in an order determined by the queuing means.
- 30 9. Apparatus according to Claim 8 wherein the queuing means comprises sorting means for sorting received service requests according to service type and allocating each received service request to a service request queue appropriate to its service type.
10. Apparatus according to Claim 9 wherein at least one service type is associated with an urgency characteristic, a first service type indicating that a

service request requires immediate delivery and a second service type indicating that a service request does not require immediate delivery.

11. Apparatus according to Claim 10 wherein at least one resource for use in
5 providing communications services is pre-allocated for use in providing a service in response to service requests of the first service type.

12. Apparatus according to any one of the preceding claims wherein the interpretation means comprises one or more interpretation tables.

10

13. Apparatus according to any one of the preceding claims wherein service requests received at the connection channel interface comprise a telephone number for use by the interpretation means in relating a service request to a computer-based application program.

15

14. Apparatus according to Claim 13, said telephone number being specific to the connection channel by means of which the service request is delivered to the interface, the interpretation means being adapted to relate the service request to a computer-based application program in accordance with the connection channel by 20 means of which it was delivered.

15. Apparatus according to any one of the preceding claims wherein the interpretation means comprises an editable computer file.

25 16. Apparatus according to any one of the preceding claims wherein the interpretation means is updatable while the platform is in use.

17. Apparatus according to any one of the preceding claims wherein the interpretation means comprises an editable computer file stored on hard disc.

30

18. Apparatus according to any one of the preceding claims which further comprises resource allocation means for allocating resources to a computer-based application program for use in providing a service.

19. Apparatus according to claim 18 wherein said resource allocation means is adapted to allocate at least some resource to a computer-based application program in advance of that application program being run.
- 5 20. Apparatus according to either one of claims 18 and 19 wherein said resource allocation means is adapted to allocate at least some resource to a computer-based application program subsequent to initiation of that application program being run.
- 10 21. Apparatus according to any one of the preceding claims wherein service requests received at the data network interface conform to an electronic mail message format.
- 15 22. Apparatus according to Claim 21 wherein service requests received at the data network interface comprise an identifier for a service together with data for use in running the service.
- 20 23. Apparatus according to Claim 22 wherein service requests received at the data network interface comprise at least first and second data fields for said data for use in running the service.
24. Apparatus according to Claim 23 which is adapted to respond in different respective ways in accordance with data being present in said first and second data fields.
- 25 25. Apparatus according to either one of Claims 23 and 24, the apparatus further comprising parsing means for use in parsing unstructured data present in a data field of a received service request.
- 30 26. Apparatus according to any one of claims 23, 24 and 25 wherein said first data field is adapted to carry an identifier for a called party in respect of a service to be provided and said second data field is adapted to carry a plurality of identifiers for called parties in respect of a service to be provided.

27. Apparatus according to any one of claims 23, 24, 25 and 26 wherein service requests received at the data network interface further comprise a third data field for carrying message content for delivery by a service to any called party for which there is an identifier in said first or second data fields.

5

28. Apparatus according to any one of the preceding claims which further comprises a traffic management system which monitors event-related data in a communication path associated with the apparatus and is provided with processing capability to process the event-related data to provide substantially real-time traffic

10 statistics.

29. Apparatus according to Claim 28 wherein the traffic management system is arranged in use to monitor all communications received by the apparatus and to detect event-related data therein.

15

30. Apparatus according to either one of claims 28 and 29 wherein the traffic management system comprises a temporary data store which carries said event-related data for monitoring.

20 31. Apparatus according to Claim 30 wherein said temporary data store comprises a cyclic buffer arrangement.

32. Apparatus according to any one of Claims 28, 29, 30 and 31 wherein the traffic management system is provided with means to access the interpretation
25 means.

33. A method of providing a communications service which comprises receiving a service request at a computing platform, identifying a computer-based application program for use in providing a service in response to the request, and
30 running the identified application program, wherein said platform has an interface for providing access to two or more connection channels of a network for carrying voice signals and said service request is received at said interface.

34. A method according to Claim 33 wherein said identification is carried out by reference to the connection channel by means of which said service request is received at the interface.

5 35. A method of providing a communications service which comprises receiving a service request at a computing platform, identifying a computer-based application program for use in providing a service in response to the request, and running the identified application program, wherein said platform has a voice channel interface for connecting the platform to at least two connection channels
10 of a network for carrying voice signals, and said platform has a data network interface for connecting the platform to a network for carrying data, wherein said service request is received at the data network interface, said identification being carried out by reference to the service request content.

15

20

25

30





Fig.2.






Fig.4.




Fig.5.




Fig.6.




Fig. 7.




Fig.8.




Fig.9.





Fig.10.



INTERNATIONAL SEARCH REPORT

International Application No

PCT/GB 97/02608

A. CLASSIFICATION OF SUBJECT MATTER
 IPC 6 H04M3/50 H04L12/58 H04M3/42

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
 IPC 6 H04M H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 91 17616 A (UNISYS CORP) 14 November 1991	33
Y	see abstract	1-3,5-7
Y	EP 0 631 419 A (VMX INC) 28 December 1994 see abstract	1-3,5-7
P,X	US 5 633 916 A (GOLDHAGEN BRUCE ET AL) 27 May 1997 see the whole document	1-35
A	WO 93 18610 A (UNISYS CORP) 16 September 1993	
A	US 5 619 555 A (FENTON WAYNE ET AL) 8 April 1997	
A	EP 0 650 284 A (AT & T CORP) 26 April 1995	
		-/-

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"Z" document member of the same patent family

1

Date of the actual completion of the international search

1 December 1997

Date of mailing of the international search report

17/12/1997

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Vandevenne, M

INTERNATIONAL SEARCH REPORT

International Application No

PCT/GB 97/02608

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 92 14314 A (SIMPACT ASSOCIATES INC) 20 August 1992	
A	EP 0 565 850 A (IBM) 20 October 1993 -----	

INTERNATIONAL SEARCH REPORT

Information on patent family members

Int. Appl. No.

PCT/GB 97/02608

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 9117616 A	14-11-91	US 5133004 A AU 7906891 A CA 2081574 A DE 69120811 D DE 69120811 T EP 0527950 A US 5384829 A	21-07-92 27-11-91 08-11-91 14-08-96 06-03-97 24-02-93 24-01-95
EP 0631419 A	28-12-94	AT 136183 T CA 2115822 A DE 69400115 D DE 69400115 T US 5557659 A	15-04-96 23-12-94 02-05-96 14-11-96 17-09-96
US 5633916 A	27-05-97	NONE	
WO 9318610 A	16-09-93	US 5323450 A CA 2130112 A EP 0629329 A JP 7504553 T	21-06-94 16-09-93 21-12-94 18-05-95
US 5619555 A	08-04-97	NONE	
EP 0650284 A	26-04-95	US 5475738 A CA 2132360 A JP 7183961 A	12-12-95 22-04-95 21-07-95
WO 9214314 A	20-08-92	US 5127003 A US 5351276 A EP 0524297 A	30-06-92 27-09-94 27-01-93
EP 0565850 A	20-10-93	US 5333266 A JP 2113103 C JP 6044157 A JP 8027766 B	26-07-94 21-11-96 18-02-94 21-03-96